

AD-A159 693


TOWARD A SURFACE SKETCH(U) MASSACHUSETTS INST OF TECH
CAMBRIDGE ARTIFICIAL INTELLIGENCE LAB J PONCE ET AL.
APR 85 AI-M-824 N00014-80-C-0505


1/1

UNCLASSIFIED

F/G 12/1

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

(2)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AIM 824	2. GOVT ACCESSION NO. AD-A159 693	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Toward a Surface Sketch		5. TYPE OF REPORT & PERIOD COVERED AI memo
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Jean Ponce & Michael Brady		8. CONTRACT OR GRANT NUMBER(s) N00014-80-C-0505 N00014-77-C-0389
9. PERFORMING ORGANIZATION NAME AND ADDRESS Artificial Intelligence Lab. 545 Technology Sq. Cambridge, MA 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Blvd. Arlington, VA 22209		12. REPORT DATE April, 1985
		13. NUMBER OF PAGES 30
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Arlington, VA 22217		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Distribution is Unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES None		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Vision Edge Finding 3D Vision Robotics Surface Representation		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This paper reports progress toward the development of a representation of significant surface changes in dense depth maps. We call the representation the SURFACE PRIMAL SKETCH by analogy with representations of intensity changes, image structure, and changes in curvature of planar curves. We describe an implemented program that detects, localizes, and symbolically describes: STEPS, where the surface height function is discontinuous; ROOFS, where the surface is continuous but the surface normal is discontinuous; SMOOTH JOINS, where the surface normal is continuous but a principal curvature is dis-		

DTIC
ELECTE
OCT 8 1985
A

AD-A159 693

DTIC FILE COPY

DD FORM 1473
1 JAN 73EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20)

continuous and changes sign; and SHOULDERS, which consist of two roofs and correspond to a STEP viewed obliquely. We illustrate the performance of the program on range maps of objects of varying complexity.

(A)

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY


A.I. Memo 824

April, 1985



Toward a Surface Primal Sketch

Jean Ponce
Michael Brady



Abstract: This paper reports progress toward the development of a representation of significant surface changes in dense depth maps. We call the representation the *Surface Primal Sketch* by analogy with representations of intensity changes, image structure, and changes in curvature of planar curves. We describe an implemented program that detects, localizes, and symbolically describes: *steps*, where the surface height function is discontinuous; *roofs*, where the surface is continuous but the surface normal is discontinuous; *smooth joins*, where the surface normal is continuous but a principal curvature is discontinuous and changes sign; and *shoulders*, which consist of two roofs and correspond to a *step* viewed obliquely. We illustrate the performance of the program on range maps of objects of varying complexity, *is illustrated*.

The authors
Additional keywords:
artificial intelligence; three dimensional vision; robotics;
images; mathematical models; algorithms; geometry.

© Massachusetts Institute of Technology, 1984

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's Artificial Intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505, the Office of Naval Research under contract number N00014-77-C-0389, and the System Development Foundation. This work was done while Jean Ponce was a visiting scientist on leave from INRIA, Paris, France.

85 10 07 038

1. Introduction

This paper describes an implemented program that detects, localizes, and symbolically describes certain types of significant surface changes in dense depth maps. Specifically, we find

- *steps*, where the surface height function is discontinuous;
- *roofs*, where the surface is continuous but the surface normal is discontinuous;
- *smooth joins*, where the surface normal is continuous but a principal curvature is discontinuous and changes sign; and
- *shoulders*, which consist of two roofs and correspond to a *step* viewed obliquely.

Figures 4, 7, 9, and 10 show the idealized instances of these surface changes that are the basis of the mathematical models used by the program. Section 6 illustrates the performance of the program on range maps of objects of varying complexity.

The work reported here continues our investigation of surface descriptions based on the concepts of differential geometry [Brady, Ponce, Yuille, and Asada, 1985]. Section 2 summarizes our ideas and shows the kind of geometric ("CAD") description we are aiming at. An important component of our work is the identification and isolation of a set of *critical surface curves*, including significant surface changes.

To this end, we report progress on the development of a representation of significant surface changes. We call the representation the *Surface Primal Sketch* by analogy with:

- Marr's [1976] *Primal Sketch* representation of significant intensity changes;
- Asada and Brady's [1984] *Curvature Primal Sketch* representation of significant curvature changes along planar contours; and
- Haralick, Watson, and Laffey's [1983] *Topographic Primal Sketch* representation of image structure.

In each case, there are three distinct problems: (i) to *detect* significant changes; (ii) to *localize* those changes as accurately as possible; and (iii) to *symbolically describe* those changes. We follow the approach of Asada and Brady [1984], as sketched in Section 3. A key component of that approach is scale space filtering, pioneered by Witkin [1983]. Yuille and Poggio [1983a, 1983b] have proved that, in principle, scale space filtering enables a discontinuity to be accurately localized. Canny [1983] uses the smallest scale at which a given intensity change can be detected to most accurately localize it.

Brady, Ponce, Yuille, and Asada [1985] report initial experiments that adapt Asada and Brady's [1984] algorithm to find surface changes. In Section 3, we describe a number of problems, both mathematical and implementational, with that approach. Section 4 describes a robust algorithm that solves the problems enumerated in Section 3 to find *roofs*, *steps*, *smooth joins*, and *shoulders*. *Roofs* are found from extrema of curvature (positive maxima and negative minima), whereas *steps*, *shoulders*, and *smooth joins* are found from parabolic points: zero crossings of the Gaussian curvature. We use scale-space behavior to discriminate *steps*, *shoulders*, and *smooth joins*. Section 6 shows the algorithm at work.

2. Background

In this section, we recall some of the main features of our work on representing visible surfaces. We work with dense depth maps that are the output of "shape-from" processes such as stereo or, more usually, direct ranging systems. There are three principal problems to be addressed:

1. Finding surface intersections. These enable the *description* of the depth map to be partitioned into a set of smooth surface patch descriptions. This is the problem addressed in the present paper. Surface intersections do not, in general, partition the depth map. Consider, for example, a bulbous end of an American telephone handset (Figure 1). The surface intersection marked on the figure peters out by the time the cylindrical portion is reached. Each surface intersection has an associated description that includes its type (*step*, *roof*, *smooth join*, etc.). In general, the type of surface intersection may vary along its length [Huffman 1971, Turner 1974]. If a surface intersection has a special property, such as being planar, that property is included in the description.
2. Generating descriptions for the smooth surface patches that result from the partitioning in (1). This is the problem addressed by Brady, Ponce, Yuille, and Asada [1985], who introduce a representation called *Intrinsic Patches*. This is discussed further below.
3. Matching surface descriptions to a database of object models that integrate multiple viewpoints of a surface. We have not yet addressed this problem. Grimson and Lozano-Pérez [1984], Faugeras, Hebert, Pauchon, and Ponce [1984], and Faugeras and Hebert [1983, 1985] have made a solid start on the problem, though they restrict attention to the case of polyhedral approximations to surfaces. However, Faugeras and Hebert [1985] illustrate the advantages of representations based on sculptured surfaces. Bou [1984], Little [1985], and Ikeuchi and Horn [1984] have developed the Extended Gaussian Image (EGI) representation for recognition and attitude determination. The EGI is an information-preserving representation only for complete maps of convex objects, a rare situation in practice. Not much has been done to extend the representation to handle non-convex objects.

The *Intrinsic Patch* representation that we are developing is based on concepts of differential geometry, principally because it provides a hierarchy of increasingly stringent surface descriptions. A surface may simply be (doubly) curved, but, in some cases, it may be ruled, even developable, even conical. Our aim is to find the most appropriate and most stringent descriptors for portions of a surface. If, for example, there is a connected region of umbilic points, indicating that part of the surface is spherical, then it is made explicit, as is the center of the corresponding sphere (Figure 2). If there is a portion of the surface that is determined to be part of a surface of revolution, it is described as such, and the axis is determined (see Figures 2 and 16).

Similarly, if there is a line of curvature or an asymptote that is planar or whose associated curvature (principal curvature or geodesic curvature respectively) is constant, then it is made explicit. For example, the asymptote (which in this particular case is also a parabolic line) that marks the smooth join of the bulb and the stem of the lightbulb in Figure 2, as well as the surface intersections marked on the oil bottle in Figure 16,




Figure 1. A telephone handset illustrates that surface intersections on curved surfaces do not, in general, partition the surface into a patchwork of smooth components.

are noted in the representation. The program described in Section 3 cannot compute the asymptote on the lightbulb; but that described in Section 4 can. We may associate a description with a curve that is a surface intersection; but only if it has an important property such as being planar. For example, a slice of a cylinder taken oblique to the axis of the cylinder produces a planar curve of intersection. Machining operations such as filleting tend to produce planar curves. Similarly, the intersection of a finger of a dextrous robot hand [Salisbury and Craig 1982, Jacobsen et. al. 1984, 1985] and an object surface is planar. On the other hand, the intersection of two cylinders is not a planar curve.

Figure 2b illustrates the representation we are aiming at. The stem of the lightbulb is determined to be cylindrical, because it is ruled and because it is a surface of revolution. We can compute the axis of the stem. The bulb is determined to be a portion of a sphere, because it is a connected region of umbilic points. The center of the sphere can be computed. Similarly, the center of the spherical portion that forms the threaded end can be determined. The stem is smoothly joined to the bulb. Moreover, the axis of the cylindrical stem passes through the centers of the spheres defined by the bulb and threaded end. This distinguishes the diameters of each sphere that are collinear with the stem axis, showing that the lightbulb is a surface of revolution. All of Figure 2b can be computed by the algorithms described in this paper and in Brady, Ponce, Yuille, and Asada [1985], except for the rightmost column, which relates to the inferences that derive from attaching the spherical portions to the cylindrical stem. Currently, we are working on the inference engine (see also Kapur, Mundy, Musser, and Narendran [1985]).

3. Surface intersections from lines of curvature

Asada and Brady [1984] introduce a representation, called the *Curvature Primal Sketch*,





Figure 2. The representation of a lightbulb. a. The dotted region consists of umbilic points, indicating that the bulb is spherical. The parallel lines are the meridians of the cylindrical stem. The parallels, which are also rulings, are not shown. b. The representation that we are working towards for the lightbulb. All save the rightmost column can be automatically computed by existing programs.

of the significant changes of curvature along a planar curve. We review that work here because our extension to surfaces follows an analogous development. Asada and Brady describe an algorithm that not only detects and localizes significant changes, but describes those changes symbolically. The simplest descriptor is *corner*, where two arcs meet continuously but where the tangent is discontinuous. Other descriptors are composed of two or more instances of the corner model. The curve that is input to the algorithm is represented by its tangent $\theta(s)$, where s is the intrinsic arclength coordinate. The algorithm is based on a mathematical analysis of a set of models that are idealized instances of the descriptors. For example, the *corner* model is formed by the intersection of two circles. Note that this is intended as a *local* approximation to a corner to facilitate analysis. It does not prejudice the subsequent approximation of the contour to be piecewise circular. Rather it suggests a set of knot points for any appropriate spline approximation.

Asada and Brady derive a number of salient features of the curvature of the models as they vary with the scale of the smoothing (Gaussian) filter. For example, a corner generates a curvature maximum, equivalently a positive maximum flanking a negative minimum in the first derivative of curvature. The height and separation of these peaks varies in a characteristic fashion over scale. The salient features are the basis of the tree matching algorithm that locates a curvature change and assigns it a descriptor. Note

that the distance between peaks varies approximately linearly (in arclength) with scale.

In the next section we develop an analogous mathematical framework for significant surface changes. Our surface analysis is local and based upon smoothing (locally) cylindrical functions with a Gaussian distribution. This is because of the following Theorem, proved in Brady, Ponce, Yuille, and Asada [1985].

The Line of Curvature Theorem: The convolution of a cylindrical surface with a Gaussian distribution is cylindrical. In more detail, let $f(x, y, z)$ be a surface that is the cross product of a planar curve and a straight line. The lines of curvature of the convolution of f with a Gaussian distribution are in the plane of the curve and parallel to the generating line.

In vector notation, a cylindrical surface has the form $r(x, y) = xi + yj + f(x)k$, and consists of parallel instances of a curve $f(x)$ in the $x - z$ plane. Our models for *roof*, *step*, *smooth join*, and *shoulder* correspond to different choices for the function $z = f(x)$.

The curvature of the smoothed curve is given by the non-linear expression

$$\kappa_{smooth}(x) = \frac{z''_{smooth}}{(1 + z'^2_{smooth})^{\frac{3}{2}}}. \quad (1)$$

Since Asada and Brady [1984] could work with tangent directions $\theta(s)$ along a planar curve, the curvature was the linear expression $d\theta(s)/ds$, so that the curvature of a smoothed contour is simply equal to the smoothed curvature of the original contour. This is *not* the case for surfaces represented as height functions $z(x, y)$. For example, the (constant) curvature of the parallels of a surface of revolution are modified (see Figure 15a). The non-linearity of curvature complicates considerably the analysis of surface change models presented in Section 4 relative to those used by Asada and Brady [1984]. Non-linearity affects smoothing too, as we discuss in Section 5.

Brady, Ponce, Yuille, and Asada [1985] used the *Line of Curvature Theorem* directly in a *two-step process* to detect, localize, and symbolically describe surface intersections, as follows:

1. Compute the lines of curvature on the surface;
2. Compute significant changes of curvature along the lines of curvature found in the first step.

The lines of curvature are computed using a best-first region growing algorithm [Brady, Ponce, Yuille, and Asada 1985]. A *good continuation* function is defined between neighboring points of the surface. The function involves the Cartesian distance between the points and the inner product of the tangent vectors corresponding to the curvature principal directions at the two points. The region growing algorithm joins the point pair whose good continuation function is globally maximum, and incorporates the new link into the developing set of lines of curvature. Brady, Ponce, Yuille, and Asada [1985] show several illustrations of the algorithm's performance. In the second step of finding surface intersections, Asada and Brady's algorithm for computing the Curvature Primal Sketch, described in the previous section, is applied to the lines of curvature in turn.

The two step process has been tested on the objects shown in Brady, Ponce, Yuille, and Asada [1985]: a lightbulb, a styrofoam cup, and a telephone receiver. It is robust

and gives good results, suggesting that the method has competence. Nevertheless, there are several problems with the method:

- **The method is inefficient.** Typical running times on a lisp machine for a smoothed depth map that is 128 points square are of the order of one hour. Much of the time is spent on further smoothing each of the (typically hundreds of) lines of curvature at multiple scales, as required by the Curvature Primal Sketch algorithm.
- **Multiple multiple smoothing is mathematically confused.** The raw surface data is smoothed at multiple scales σ_i , giving a set of surfaces z_i . The Curvature Primal Sketch algorithm further smooths the lines of curvature of z_i at multiple scales σ_j yielding a set of smoothed lines of curvature $r_{ij}(s_{ij})$. There is no obvious relation between the scales σ_i and σ_j .
- **Discretization makes implementation difficult.** The lines of curvature of an analytic surface form a dense orthogonal web. The (smoothed) depth maps we work with are discrete approximations to analytic surfaces. In practice, the lines of curvature found by the two step process are sometimes broken. The lines of curvature near the perceptual join of the stem and bulb of the lightbulb shown in Figure 2 illustrates this problem. This is due in part to quantisation effects, but is also because the principal directions change rapidly near surface discontinuities. This is why the *smooth join* between the bulb and the stem of the lightbulb is not found by the two step process.
- **The *Line of Curvature Theorem* only applies locally.** In practice, few surfaces are cylindrical in the sense of the *Line of Curvature Theorem*. The Theorem is only approximately true in general, and then only locally. The application of the Curvature Primal Sketch algorithm in the second step does not respect this.
- **Lines of curvature on smoothed surfaces are not planar curves.** The models that are embodied in the Curvature Primal Sketch algorithm are not a complete set for surface intersections.

The success of the two step process suggests that the method is on the right track. The problems just enumerated suggest that reducing the problem to apply an existing algorithm developed for planar curves, though expedient, is wrong. Together, these observations suggest that a real two-dimensional extension of the Curvature Primal Sketch should be developed. The next section reports our progress toward such an extension.

4. Toward a surface primal sketch

4.1. A three-step process

In this section we develop a method for finding certain types of changes in the height of a surface that overcomes the difficulties of the two-step process described in the previous section. The types of changes we have analyzed and implemented are as follows: *steps*, where the surface height function is discontinuous; *roofs*, where the surface is continuous but the surface normal is discontinuous; *smooth joins*, where the surface normal is

continuous but a principal curvature is discontinuous and changes sign; and *shoulders*, which consist of two roofs and correspond to a *step* viewed obliquely. It turns out that *roofs* consist of extrema of the dominant curvature: that is, maxima of the positive maximum curvature or minima of the negative minimum curvature. On the other hand, *steps*, *smooth joins*, and *shoulders* consist of parabolic points, that is zero crossings of the Gaussian curvature. They are distinguished by their scale space behavior.

We have implemented the following *three-step process* (Figure 3) that is illustrated in the examples presented in Section 6:

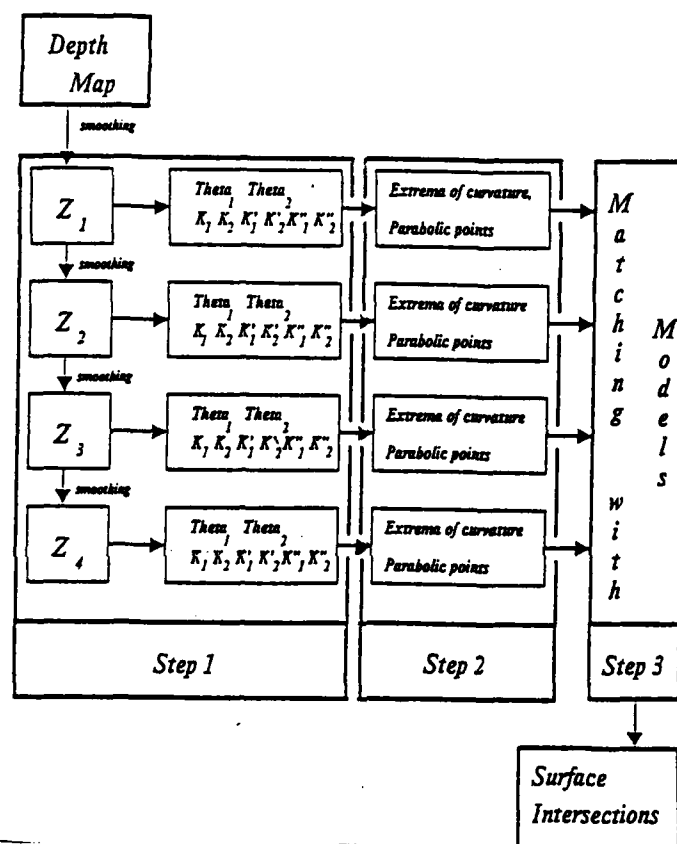


Figure 3. Schematic of the three-step process described in this paper and in Brady, Ponce, Yuille, and Asada [1985]. The analysis of the set of models and the matching algorithm are the topic of the present paper. Results of running the process are shown in Section 6.

- 1 Smooth the surface with Gaussian distributions at a set of scales σ_i , yielding surfaces z_i . Compute the principal directions and curvatures everywhere;
- 2 In each smoothed surface z_i , mark the zero-crossings of the Gaussian curvature and the (directional) extrema of the dominant curvatures;
- 3 Match the descriptions of the surfaces z_i to find points that lie on *roof*, *step*, *smooth join*, and *shoulder* surface discontinuities.

Note that Brady, Ponce, Yuille, and Asada [1985] investigated parabolic lines and lines of curvature as *global* descriptors of surfaces. They suggest that such a line needs additional global properties, such as planarity, to be perceptually important. In this paper, we are

interested in parabolic points and curvature extrema as *local cues* for significant surface intersections.

We discuss smoothing in more detail in the next Section. The computation of principal curvatures is described in Brady, Ponce, Yuille, and Asada [1985]. The next four subsections analyze *steps*, *roofs*, *smooth joins*, and *shoulders*. Subsequent subsections discuss the matching algorithm and further work that is needed to elaborate the model set.

Is it necessary to use multiple scales to find surface intersections? Arguments supporting multiple scales for edge finding in images have been advanced elsewhere [see Marr and Hildreth 1980, Canny 1983, Witkin 1983]. However, it might be supposed that it would be sufficient to smooth depth maps with a single coarse filter. Figures 12b and 15a show that this is not so. Even after thresholding, there is still a large curvature extremum in the neck of the bottle running parallel to the axis. This extremum is an artefact of non-linear smoothing, and it cannot be eliminated at a single scale. Instead, we reject it because it does not change over scale in the characteristic manner of a *roof*.

4.2. Step discontinuities

A *step* occurs when the surface itself is discontinuous. The model we use consists of two slanted half planes whose normals lie in the $x - z$ plane. They are separated by a height h at the origin (Figure 4). Using the line of curvature theorem, we study the one dimensional formulation of this model.




Figure 4. The *step* model, consisting of two slanted planes separated by a height h at the origin. The *roof* model corresponds to the case $h = 0$ and $k_1 \neq k_2$.

Let the curve $z = f(x)$ be defined by

$$z = \begin{cases} k_1 x + c, & x < 0; \\ k_2 x + c + h, & x > 0; \end{cases} \quad (2)$$

where h and c are constants. In this expression, h is the height of the *step*. We now derive the result of smoothing this function with a Gaussian distribution at a given scale σ . To obtain a symmetric form for this smoothed version, we introduce the two following parameters:

$$k = (k_1 + k_2) / 2$$

$$\delta = k_2 - k_1$$

If we denote the smoothed curve ($G_\sigma * z$) by z_σ , we then obtain

$$z_\sigma = c + \frac{h}{\sigma\sqrt{2\pi}} \int_{-\infty}^x \exp(-\frac{t^2}{2\sigma^2}) dt$$

$$+ kx + \frac{\delta x}{\sigma\sqrt{2\pi}} \int_0^x \exp(-\frac{t^2}{2\sigma^2}) dt \quad (3)$$

$$+ \frac{\delta\sigma}{\sqrt{2\pi}} \exp(-\frac{x^2}{2\sigma^2}),$$

The first and second derivatives of z_σ are given by

$$z'_\sigma = k + \frac{\delta}{\sigma\sqrt{2\pi}} \int_0^x \exp(-\frac{t^2}{2\sigma^2}) dt + \frac{h}{\sigma\sqrt{2\pi}} \exp(-\frac{x^2}{2\sigma^2}); \quad (4)$$

$$z''_\sigma = \frac{1}{\sigma\sqrt{2\pi}} (\delta - \frac{hx}{\sigma^2}) \exp(-\frac{x^2}{2\sigma^2}) \quad (5)$$

In particular, the curvature κ , given by Equation (1), has a zero crossing at the point $x_\sigma = \sigma^2 \delta / h$. This is at the origin if and only if $k_1 = k_2$, otherwise, the distance from x_σ to the origin is proportional to σ^2 . This is illustrated in Figure 5 for the *step* between the cylindrical body and the cylindrical base of the oil bottle shown in Figure 16. From Figure 5, we calculate δ/h to be 0.105. The actual height of the step is about 1.5 millimeters. By the way, the position of the zero crossing shown in Figure 5 moves by about 3 pixels over one octave.

Using the fact that the second derivative of z_σ is zero at x_σ , it is easy to show that

$$\frac{\kappa''}{\kappa'}(x_\sigma) = \frac{z''''_\sigma}{z'''_\sigma}(x_\sigma) = -\frac{2\delta}{h} \quad (6)$$

So the ratio of the second and first derivatives of the curvature at the zero crossing is constant over the scales. Calculating δ/h this way gives 0.11, which is close to the value given by the slope in Figure 5. This suggests that one ought not be overly coy about computing first and second derivatives of curvature of appropriately smoothed versions of a surface, even though they correspond to third and fourth derivatives.

It is possible to substantially reduce the effect of this problem by smoothing in intrinsic coordinates. At each point of the surface, the normal is estimated. Instead of smoothing z , the surface point is moved along its normal a distance that depends upon the projected distances of the points neighbors from the tangent plane. In the case that the normal faces the viewer, this is equivalent to the previous technique. However, the result is no longer orientation-dependent. Figure 15b illustrates the computation of curvature after smoothing the oil bottle by this method. The drawback with the technique is the computation it requires one to compute the tangent plane at every point.

The technique described in Brady, Ponce, Yuille, and Asada [1985] has been used in all the examples presented here, as it represents a good tradeoff between computational efficiency and faithful rendering of the smoothed surface.

6. Examples

In this section, we present a number of examples of the surface discontinuities found on simple objects by our algorithm. In all the examples, we use four different scales corresponding to 20, 40, 60, and 80 iterations of the smoothing filter described in Brady, Ponce, Yuille, and Asada [1985]. Viewing the resulting centrally-limiting Gaussian distributions as approximately bandpass filters, they span one octave.

Figure 16 shows the final output of the algorithm for the oil bottle. The points detected during the matching step are linked together using a connected components exploration algorithm. The smallest components (less than 3 or 4 pixels) are then removed. Conversely, points may have been missed during the previous phases, creating gaps in the lines that are found. These gaps are filled by adding points that have characteristics compatible with the detected points. The bottle is finally segmented into six parts, separated by three step edges and two roofs.

Brady, Ponce, Yuille, and Asada [1985, Figures 18 and 19] showed that the coffee cup shown in Figure 17 is best represented as the join of a cylindrical body and a tube surface that corresponds to the handle. Here we show that the handle can be separated from the body using the algorithms described in this paper. Note that the surface intersections are of type *roof*.

The third example shows the surface intersections found on a telephone handset, Figure 18. All the major intersections have been found. The representation is not symmetric because the handset was not quite perpendicular to the scanner, causing part of the surface to be occluded. Note that the surface intersections are more reliably detected at the coarsest scale, but are more accurately localized at the finest scale.

The surface intersections found on a few simple tools, namely a hammer, a drill, and the head of a screwdriver are showed in Figure 19.

Figure 20 shows the surface intersections found on an automobile part that has featured in several papers by the group at INRIA. On this complicated object, global lines of curvature have no signification, so the Curvature Primal Sketch would not perform well. Notice in particular the circular step edge found on the left "head" of the part: it corresponds to a shallow depression whose depth is about one millimeter. This is approximately at the resolution limit of the laser scanner, and underlines the practical significance of the algorithms described here.




Figure 14. a. Raw data from a cross section of an oil bottle after scanning using the INRIA system. b. Smoothing across surface boundaries with a Gaussian mask that is applied everywhere. c. Gaussian smoothing using repeated averaging and computational molecules. (Reproduced from Brady, Ponce, Yuille, and Asada [1985, Figure 12])




Figure 15. a. The curvature computed from the smoothed surface shown in Figure 14c. Small orientation-dependent errors in smoothing are magnified. The first figure is the neck, the second part of the body. b. The curvature on the slices shown in a. computed using the intrinsic coordinate method described in the text

decreases at a certain distance from the boundary. This creates an inflection point, which in turn creates an extremum of curvature.

types of intersections, and be associated to an ancestor of each of these types. The use of multiple scales is usually sufficient to disambiguate between the different cases. If several interpretations remain after the finest level has been taken into account, the one with the best cumulated compatibility score is chosen.

Note that, although we have analyzed the behavior of the position of the characteristic point of each of our models, and have found it simple and reliable, we do not use it in the compatibility functions. The reason is that, for each intersection found, the real origin on the surface is unknown. This implies that at least three matched points are necessary to estimate the parameters of the motion of the extremum or zero crossing across the scales, and makes the measure of this movement unsuitable for a scale-to-scale tracking. However, the estimation of the movement parameters could be used for an *a posteriori* verification of the surface intersections found.

5. Smoothing a surface with a Gaussian distribution

Brady, Ponce, Yuille, and Asada [1985] discuss techniques for smoothing a depth map with a Gaussian distribution. The main difficulty stems from bounding contours, where the surface normal turns smoothly away from the viewer, and where there is typically a substantial depth change between points on the surface of the object and the background. In general, the bounding contour is easy to find, even with a simple edge operator or by thresholding depth values. The problem is how to take the boundary into account when smoothing the surface.

Brady, Ponce, Yuille, and Asada [1985] observed that if the smoothing filter is applied everywhere, the surface "melts" into the background and changes substantially. Figure 14 is reproduced from Brady, Ponce, Yuille, and Asada [1985, Figure 13c] and shows this. They suggested instead using repeated averaging [Burt 1981, 1983] as well as adapting Terzopoulos' [1983] technique of computational molecules to prevent leakage across depth boundaries. This smooths the surface without substantially altering it (see Figure 14c).

Here we point out a slight difficulty in smoothing surfaces using the technique illustrated in Figure 14c, and suggest a refinement. Although the smoothed surface appears to be close to the original, small orientation-dependent errors are introduced. These errors are magnified in computing the curvature (Figure 15a), to produce "false" curvature extrema near the boundary (compare the overshoot phenomenon in Terzopoulos' [1983] work on detecting surface discontinuities). The overshoots do not exemplify Gibbs' ringing as we originally thought. Instead, the phenomenon is caused by two effects:

- **The coordinate frame is not intrinsic.** The smoothing filter is applied in the $x - y$ plane, and since this is not intrinsic to the surface, the result is orientation-dependent. For example, the difference between a cylinder and its smoothed version monotonically increases to the boundary from a value of zero where the normal faces the viewer.
- **Points near the boundary don't get smoothed as much.** Such points are relatively unsmoothed as several of their computational molecules are continually inhibited. The result is that the difference between the smoothed and original surfaces




Figure 13. Matching feature points across scales. See text for details.

- **Parabolic points** are generated by several models.

Curvature extrema are only generated by *roofs* in our model set. However, parabolic points (zero crossings) may correspond to different types of discontinuities. Worse, a point may simultaneously be a zero crossing of the Gaussian curvature, and an extremum of a principal curvature! Again, our solution is to define compatibility functions for *steps*, *shoulders*, and *smooth joins*, that take into account their mathematical models. The compatibility functions are based on the behavior of the ratio of the second and first derivatives of curvature. At each scale, a point may be a candidate for several different

extremum) as it is tracked across scales. Points at the finest scale that have no ancestor at the coarsest scale on their path are eliminated. This constitutes the *detection* of surface changes. To *localize* surface changes, we note that in each case analyzed in Sections 4.2 through 4.6 the distance of the feature point from the origin increases with σ . For this reason, the position of the surface change is determined from the finest scale.

Finally, each path (the scale-space tracking of a feature point) is assigned a *local signature*: *roof*, *step*, *smooth join*, or *shoulder*, depending on the behavior of its curvature and its first and second derivatives of its curvature across the scales. That is, the path is analyzed according to the models developed in Sections 4.2 through 4.6. The parameterized local signature provides a *symbolic description* of the surface change on which the feature point lies.

Figure 13 shows the matching algorithm at work on several image slices of the oil bottle shown in Figure 12. The final result is shown in Figure 16. The three parts of Figure 13 show, for consecutive pairs of scales, points that have been matched in the same set of eighteen image slices of the oil bottle. The upper group of eighteen graphs corresponds to the coarsest scale (called "80" because it corresponds to eighty iterations of the smoothing computational molecules [Brady, Ponce, Yuille, and Asada, 1985]) being matched to the next-to-coarsest scale "60". The middle group corresponds to matching between scales 60 and 40; the bottom group to scales 40 and 20 (which is the finest scale). In a given position in the blocks of eighteen graphs, say the fourth from the left in the middle row, an image slice is tracked across the three pairs of scales. Let us consider one of the pairs of scales, say 60 and 40 shown in the middle block. Feature points that are matched are linked by a vertical line.

Matching is not straightforward:

- **Surface changes lie on space curves.**

Consider, for example, a *roof* discontinuity whose local signature is a curvature extremum. In the *Curvature Primal Sketch*, curvature extrema are isolated points along a one-dimensional curve, and this makes the construction of the trees of corresponding feature points relatively simple (see the figures in [Asada and Brady 1984]). In three dimensions, however, surface changes constitute continuous space curves. The association of an ancestor (respectively descendant) with a given feature point is often ambiguous, and this complicates the construction of the forest. This, in turn, makes difficult an *a posteriori* interpretation of this forest.

Our solution is to compute a compatibility between each matched pair of marked points. The compatibility function involves the Cartesian distance between the points and the angle between their associated principal directions, but also takes into account the *roof* model by comparing the ratio of the points curvatures to the inverse of the ratio of the associated scales. At each scale, and for each thresholded extremum, we look for an ancestor inside a square window of the previous scale image. If an ancestor with a sufficiently high score is found, then the point is kept as a potential ancestor for the next scale. Otherwise it is removed. This way, the forest is never explicitly built, and the interpretation is done during the tracking itself, as the only extrema tracked are those which correspond to potential *roofs*. In particular, this is how the artefacts due to non-linear smoothing are removed.

- curvature extrema whose values are less than a preset threshold are removed;
- zero crossings whose slopes are less than a (different) threshold are removed.

Figure 12b shows the result of thresholding the feature point sets in 12a. The thresholds vary according to scale. For example, the extremum threshold varies proportionally to $1/\sigma$, as suggested by our analysis of the *roof* model (see Section 4.3). To date, we have not derived an analogous formula for the scale space variation of the zero crossing threshold.




Figure 12. a. The extrema and zero crossings of curvature for the oil bottle at four scales that increase from left to right. The total variation in scale is one octave. b. The feature points in (a) that are above threshold. Note the curvature extrema parallel to the axis of the bottle that are artifacts of the non-linear smoothing. Note also the numerous parabolic points at the finest scale that are not thresholded. These non significant points are eliminated by the matching algorithm (Figures 13 and 16).

This is not a major problem, however, as the thresholding step is only used for selecting a set of candidates for the subsequent matching process, rather than finding the surface intersections themselves. For example, the curvature extrema parallel to the axis of the oil bottle, that are due to the non-linear smoothing (Figure 15), cannot be eliminated by thresholding, but are rejected by the matching algorithm (Figure 16) since they do not conform to a model.

The matching algorithm is a two-dimensional extension of that proposed by Asada and Brady [1984]. We track the thresholded extrema and zero crossings across scales, from coarse to fine. We obtain a forest of points, equivalent to a "fingerprint" [Yuille and Poggio 1983]. Paths in the forest correspond to a feature point (zero crossing or

we shall see, however, the analysis is quite difficult for a thin bar, even in simple cases.

We analyze a model of a *thin bar* consisting of a plateau of height h , width $2a$, and slope k_2 , resting on flat ground of slope k_1 (Figure 10). We model the *thin bar* by the function

$$z = \begin{cases} k_1 x, & x < -a; \\ h + k_2 x, & x \in [-a, +a]; \\ k_1 x, & x > +a. \end{cases} \quad (18)$$

We denote $k_2 - k_1$ by δ . The first and second derivatives of the smoothed *thin bar* are given by:

$$z'_\sigma = k_1 + \frac{\delta}{\sigma\sqrt{2\pi}} \int_{x-a}^{x+a} \exp\left(-\frac{t^2}{2\sigma^2}\right) dt + \frac{(h - \delta a)}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x+a)^2}{2\sigma^2}\right) - \frac{(h + \delta a)}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-a)^2}{2\sigma^2}\right) \quad (19)$$

$$z''_\sigma = \frac{1}{\sigma\sqrt{2\pi}} \left(\delta - \frac{(h - \delta a)(x+a)}{\sigma^2} \right) \exp\left(-\frac{(x+a)^2}{2\sigma^2}\right) - \frac{1}{\sigma\sqrt{2\pi}} \left(\delta - \frac{(h + \delta a)(x-a)}{\sigma^2} \right) \exp\left(-\frac{(x-a)^2}{2\sigma^2}\right) \quad (20)$$

These expressions simplify considerably in the case that the plateau surface is parallel to the ground, that is $\delta = 0$. In particular, in that case $z'''_\sigma = 0$. However, the curvature attains an extremum at the origin, equivalently $\kappa' = 0$, only when $k_1 = 0$. This is the case depicted in Figure 11, but it is too restrictive since it is too sensitive to changes in viewpoint. Further work is needed here. By the way, the special case $k_1 = k_2 = 0$ is that typically studied in psychophysical studies of intensity *thin bars* (eg Richter and Ullman [1982]). It would be interesting to know what is the response to *thin bars* of intensity superimposed on a linear intensity ramp.

4.7. The matching algorithm

We now use the models introduced in the previous sections to *detect* and *localize* surface intersections. We track the extrema of curvature and parabolic points found at each scale from coarse-to-fine. The tracking is directed by the particular features associated with each model. In essence, the features constitute a *local signature* of the model. This should be contrasted with the complex search for peak configurations used by Asada and Brady [1984].

We first smooth the original depth map with a Gaussian distribution at a variety of scales σ (see Figure 3). We then compute, for each smoothed version of the surface, the principal curvatures and their directions (using the method described in Brady, Ponce, Yuille, and Asada [1985]). We also compute the first and second derivatives of the principal curvatures in their associated directions. Parabolic points (zero crossings of the Gaussian curvature) are then marked, as are the (directional) maxima of the maximum curvature and minima of the minimum curvature (Figure 12a). The marked points are thresholded, according to their type:




Figure 10. Model of a *thin bar* compound surface discontinuity. It consists of a plateau of height h , width $2a$, and slope k_2 , resting on flat ground of slope k_1 .




Figure 11. Expected curvature response of a *thin bar* to Gaussian filters. At fine scales, the *thin bar* signals two separate *steps*; at coarser scales it resembles a difference-of-Gaussians. The *step* responses begin to interfere when σ equals half the separation of the risers.

we are reluctant to implement an analogous peak matching program. In practice, the peaks from a *thin bar* may cover as many as fifteen pixels, suggesting error-prone and inefficient search. In this paper, we have sought *local statements* that apply to a single zero crossing or curvature extremum and studied its scale space behavior in isolation. As

$$x_\sigma = \frac{\sigma^2}{2a} \log \left(\frac{\delta_1}{\delta_2} \right) \quad (16)$$

Using the Implicit Function Theorem as in the case of the *roof*, it is then straightforward to show that:

$$\frac{\kappa''}{\kappa'}(x_\sigma) = \frac{z_\sigma'''}{z_\sigma''}(x_\sigma) = -\frac{1}{a} \log \left(\frac{\delta_1}{\delta_2} \right), \quad (17)$$

so the ratio of the first and second derivatives at the zero crossing is constant over scales.

4.6. Thin bar and other compound discontinuities

The models considered so far involve *isolated* surface changes. Even though a *shoulder* may appear as a pair of *roofs* if it is viewed close to the normal to the riser and if the riser subtends a sufficient visual angle, its more typical behavior is like that of a *step*. As two (or more) surface changes are brought more closely together, so that the filter width σ approaches half the separation of the changes, the filter responses due to the individual changes interfere with each other. Since certain kinds of compound surface discontinuity are important for recognition and use of objects, they must be modeled and matched by the program.

This observation raises two questions: (i) which compound surface changes should be modeled and matched; and (ii) how shall instances be found by the program? Ultimately, the answer to (i) is application-dependent, though the *thin bar*, consisting of a *step* up closely followed by a *step* down, presses for inclusion (Figure 10). *Thin bars* occur as ribs on many surfaces, for example along the sides of the neck of a connecting rod. Also, it seems [Marr 1976, Richter and Ullman 1982] that the mammalian visual system is sensitive to thin intensity stripes. In the case of curvature changes along a planar curve, Asada and Brady [1984] introduce the *crank* that is analogous to a thin bar since it consists of a corner followed closely by one of opposite sign. Other compound surface changes that might be important are a rounded corner and a moulding, that is like a thin bar but with one of its risers smooth and concave. We have not studied such configurations.

Restricting attention to *thin bars* raises question (ii): how shall instances be recognised? First, let us conjecture what the curvature response to a *thin bar* might look like. We may base our conjecture on Asada and Brady's [1984] analysis of a crank, though we need to be cautious because their operators were linear. Figure 11 shows the response that might be expected, indeed the response that is provably generated in one special case (see below). Unfortunately, the response becomes substantially more complex in the general case.

Note that the *thin bar* in Figure 11 generates as many as five curvature peaks at fine scales, reducing to three at coarser scales. Note also that there appears to be a curvature peak at the origin. Asada and Brady [1984] extracted peaks at all scales and developed a matcher that linked peaks across scales. The *crank* model explicitly checked for three peaks splitting to five in the way shown in Figure 11. Matching such compound (planar curve curvature) changes was the source of the complexity of Asada and Brady's program. In view of the non-linearity of surface curvature, and the two-dimensionality of surfaces,

4.5. Shoulder discontinuities

A *step* discontinuity confounds information both about the geometry of the surface and the viewpoint. Shifting the viewpoint to the half space defined by the outward normal of the "riser" of the step typically changes the depth discontinuity to a pair of *roofs* of opposite sign whose separation again confounds geometry and viewpoint. We introduce the *shoulder* discontinuity to cater for this situation (Figure 9).




Figure 9. The *shoulder* discontinuity consists of two *roofs* of opposite sign. The shoulder appears as a *step* when viewed from the half-space defined by the inward normal of the "riser".

We may expect the scale space behavior of the *shoulder* to closely resemble that of the *step* when the projected separation $2a$ of the *roofs* is small compared to the filter size σ , perhaps becoming more like a pair of *roofs* as the viewpoint shifts. This is what we find.

We model the *shoulder* by the function

$$z = \begin{cases} k_1 x + (k_1 - m)a, & x < -a; \\ mx, & x \in [-a, +a]; \\ k_2 x + (m - k_2)a, & x > +a \end{cases} \quad (14)$$

If we denote $k_1 - m$ by δ_1 and $k_2 - m$ by δ_2 , then $\delta_i \neq 0$, and δ_2/δ_1 is positive (otherwise the curve is always convex, or always concave). It is easy to show that the second derivative of the shoulder, smoothed at scale σ is

$$z''_\sigma = \frac{\delta_2}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-a)^2}{2\sigma^2}\right) - \frac{\delta_1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x+a)^2}{2\sigma^2}\right) \quad (15)$$

Since δ_2/δ_1 is assumed positive, we deduce from Equation (15) that the curvature has a zero crossing. The location of the zero crossing is given by:

In particular, we deduce from Equation (11) that the curvature has a zero crossing if and only if

$$\frac{1}{\sqrt{2\pi}} \int_0^{z/\sigma} \exp\left(-\frac{u^2}{2}\right) du = -\frac{c}{\delta} \quad (12)$$

This equation has a solution if and only if the absolute value of c/δ is less than $1/2$, which simply corresponds to c_l and c_r having opposite signs. It follows that *smooth joins* of the sort shown in Figure 7a generate a zero crossing in curvature, hence a parabolic point on the surface. Those shown in Figure 7b do not.

For example, the parabolic lines found on the lightbulb shown in Figure 8 are *smooth joins*. The two-step process utilising the Curvature Primal Sketch algorithm, discussed in the previous section, failed to find the *smooth joins* (see Brady, Ponce, Yuille, and Asada [1985, Figure 1]).




Figure 8. Parabolic lines found by the program described in this Section for the lightbulb shown in Figure 2. The smooth join between the stem and bulb were not found by the program described in Section 4.

Equation (12) implies that the distance from the zero crossing location x_σ to the origin is proportional to σ . Using this property and the fact that z'' is zero at x_σ , it is then easy to use the Implicit Function Theorem to show that

$$\frac{\kappa''}{\kappa'}(x_\sigma) = \frac{z_\sigma''''}{z_\sigma'''}(x_\sigma) = \frac{\gamma}{\sigma}, \quad (13)$$

for some constant γ . It follows that the ratio of the second and first derivatives of the curvature in x_σ is inversely proportional to σ . This scale space behavior allows us to discriminate zero crossings due to steps from those due to smooth joins.




Figure 7. The model for a *smooth join* consists of two parabolas meeting smoothly at the origin. a. The curvature changes sign generating a parabolic point on the surface. b. The curvature does not change sign. Such *smooth joins* are typically perceivable only when there is a large, discontinuous jump in curvature.

The continuity and differentiability of the curve at $x = 0$ imply that $b_l = b_r = b$, say, and $a_l = a_r = a$, say. As in the case of the *step*, we introduce the parameters

$$c = (c_l + c_r) / 2$$

$$\delta = c_r - c_l$$

We can express the surface, smoothed at the scale σ , as

$$\begin{aligned} z_\sigma = & \frac{1}{2} \left(c + \frac{\delta}{\sigma\sqrt{2\pi}} \int_0^x \exp\left(-\frac{t^2}{2\sigma^2}\right) dt \right) x^2 \\ & + \left(b + \frac{\delta\sigma}{2\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) \right) x \\ & + \left(a + \frac{1}{2}c\sigma^2 + \frac{\delta\sigma}{2\sqrt{2\pi}} \int_0^x \exp\left(-\frac{t^2}{2\sigma^2}\right) dt \right) \end{aligned} \quad (9)$$

The first and second derivatives of z_σ are now given by:

$$\begin{aligned} z'_\sigma = & \left(c + \frac{\delta}{\sigma\sqrt{2\pi}} \int_0^x \exp\left(-\frac{t^2}{2\sigma^2}\right) dt \right) x \\ & + \left(b + \frac{\delta\sigma}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) \right) \end{aligned} \quad (10)$$

$$z''_\sigma = c + \frac{\delta}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) \quad (11)$$




Figure 6. Scale space behavior of the *roof* discontinuity between the cylindrical neck and the conical shoulder of the oil bottle shown in Figure 16. a. The position of the negative minimum of curvature varies linearly as a function of scale as predicted by the analysis. b. The curvature is directly proportional to $1/\sigma$, as predicted.

can be estimated accurately. We use non-maximum suppression [Canny 1983] to reject local extrema. The location of the peak, its height, its type (maximum or minimum), and its orientation, are the features we use for the subsequent matching over scales.

4.4. Smooth join discontinuities

In certain circumstances, one can perceive surface changes where both the surface and its normal are continuous, but where the curvature is discontinuous. We call such a surface change a *smooth join* discontinuity. If the curvature changes sign at a *smooth join*, the surface has a parabolic point. As we shall see, such changes can be found from zero crossings of a principal curvature. It is well-known (see Asada and Brady [1984] for discussion and references) that *smooth joins* where the curvature do not change sign are perceptible only when the (discontinuous) jump in curvature is "sufficiently large". In such a case, there is not a zero crossing of curvature; rather there is a level crossing, and the curvature typically inflects. We do not yet have a complete analysis of that case.

Our model of a *smooth join* consists of two parabolas that meet smoothly at the origin (the curve is differentiable). Figure 7 shows the two distinct cases of the model. (Though the two cases appear to be perceptually distinct and lead to different matching criteria, they are governed by the same Equation (8), so it is convenient to analyze them together at first.)

Consider the curve $z(x)$ defined by

$$z = \begin{cases} \frac{1}{2}c_l x^2 + b_l x + a_l, & x < 0; \\ \frac{1}{2}c_r x^2 + b_r x + a_r, & x > 0; \end{cases} \quad (8)$$




Figure 5. Variation over scale of the position of the zero crossing of the curvature of the smoothed step between the cylindrical body and the cylindrical base of the oil bottle shown in Figure 16. The abscissa is σ^2 , the ordinate the position of the zero crossing. The height of the step is about 1.5 millimeters. The slope is $\delta/h = 0.105$.

4.3. roof discontinuities

A *roof* occurs when the surface is continuous, but the surface normal is discontinuous. Specializing Equations (2) to (5) to the case $h = 0$, we obtain

$$\kappa = \frac{1}{\sigma\sqrt{2\pi}} \frac{\delta e^{-\frac{x^2}{2\sigma^2}}}{\left[1 + \left(k + \frac{\delta}{\sigma\sqrt{2\pi}} \int_0^x \exp\left(-\frac{t^2}{2\sigma^2}\right) dt\right)^2\right]^{3/2}} \quad (7)$$

$$= \frac{1}{\sigma\sqrt{2\pi}} \frac{\delta e^{-\frac{x^2}{2\sigma^2}}}{\left[1 + \left(k + \frac{\delta}{\sqrt{2\pi}} \int_0^{x/\sigma} \exp\left(-\frac{u^2}{2}\right) du\right)^2\right]^{3/2}}$$

From Equation (7), we deduce that for a roof, we have $\kappa(x, \mu\sigma) = \kappa(x/\mu, \sigma)$. In particular, this implies that the extremum value of κ is proportional to $1/\sigma$, and that its distance from the origin is proportional to σ . This is illustrated in Figure 6 for the *roof* discontinuity between the cylindrical neck and the conical shoulder of the oil bottle shown in Figure 16. Figure 6a shows the variation in the position of the negative minimum of curvature as a function of scale. Figure 6b shows that curvature is directly proportional to $1/\sigma$.

It is also easy to show that the second derivative of the curvature, κ'' , is proportional to $1/\sigma^3$. However, we do not use this property in the current implementation of the program, relying instead on the the variation of the extremum height over scale.

We look for points that are local maxima (respectively minima) of the maximum (respectively minimum) curvature in the corresponding direction. The curvature directions




Figure 16. The oil bottle is segmented into six parts. Three step and two roof intersections are found by the algorithm described in this paper. The algorithms described in Brady, Ponce, Yuille, and Asada [1985, see Figure 18] determine the lines of curvature of the parts of the oil bottle, fit circles to the parallels, and fit axes to the centers of those circles.




Figure 17. The joins of the handle to the body of the coffee mug are computed by the algorithms described in this paper. They are determined to be of type *roof*.

Figure 21 shows the surface intersections found on the head of a connecting rod. The current state of our algorithm cannot deal with the *thin bars* located on the sides of the neck of the connecting rod, but performs well for the other intersections.

The last example is the mask of a human face (Figure 22). The program finds face




Figure 18. a. The surface intersections found on a telephone handset at the coarsest and finest scales that are approximately an octave apart. The intersections are more reliably detected at the coarsest scale; they are more accurately localized at the finest scale. b. The results of matching the changes across scales.




Figure 19. The intersections found by the program on simple tools. a. a hammer. b. a drill c. a screwdriver

features as the nose, the eyes, and the mouth. This shows its ability to deal with arbitrary curved surfaces, usually not found in man-made objects.

Although our primary concern in this paper has been with the intrinsic geometry of a surface as found by a three-dimensional vision system, one might suppose that the methods described in this paper could be applied straightforwardly to extract and




Figure 20. Surface intersections found on an automobile part (see, for example, Faugeras [1982, 1984]). The circular step edge found on the left "head" of the part corresponds to a shallow depression whose depth is about one millimeter. This is approximately at the resolution limit of the laser scanner.




Figure 21. Surface intersections found on a connecting rod. Only the head of the part is shown. The inclusion of the *thin bar* models in the algorithm would allow a fine description of the neck of this object.

interpret significant intensity changes in images, considered as surfaces. To interpret intensity changes, it is necessary to take irradiance effects into account, since intensity changes do not always correspond to surface changes. Rather, they may signify reflectance or illumination changes. Extraction and interpretation was in fact the intent of Marr's




Figure 22. Surface intersections found on a mask. The mouth, the nose and the eyes are found as corners by the program.

[1976] original work on the Primal Sketch, though considerably more attention was paid to extraction than interpretation. More recently, Haralick, Watson, and Laffey [1983] have advocated a representation of image surface geometry that involves concave and convex hills, planar regions, and saddles. These geometrical aspects of the image surface are not interpreted in terms of the intrinsic geometry of the surface, or of illumination or reflectance changes.

Some preliminary work exists on interpretation of intensity changes. An early edge finder developed by Binford and Horn [Binford 1981] included filters for step, roof, and "edge effect" changes. Horn's [1977] study of intensity changes included a suggestion that occluding boundaries and reflectance changes correspond to *step* intensity changes, while concave surface intersections generate *roof* intensity changes (because of mutual illumination). Finally, Yuille [1983] suggests that certain points along lines of curvature of a surface can be extracted directly from an image. There is much scope for additional work along these lines.

Figure 23 shows an initial experiment we have carried out on applying the methods developed in this paper to image surfaces. The join of the wing to the fuselage of the airplane is determined to be *roof* changes, consistent with Horn's suggestion.

7. Acknowledgements

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's Artificial Intelligence research is provided in part by the the System Development Foundation, the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643, and the Office of Naval Research under contract number




Figure 23. Application of the methods of the paper to an intensity surface. The interest is in the type of the internal intensity changes. The join between the wings and fuselage is a roof, suggesting a concave surface intersection.

N00014-80-C-050. We thank several people who have commented on the ideas presented in this paper, particularly Ruzena Bajcsy, Bob Bolles, Jean-Daniel Boissonat, Bernard and Hilary Buxton, Olivier Faugeras, Eric Grimson, Berthold Horn, Ramesh Jain, Takeo Kanade, Christopher Longuet-Higgins, David Mumford, Joe Mundy, Tommy Poggio, Robin Stanton, Demetri Terzopoulos, and Alan Yuille. We are indebted to Olivier Faugeras and the INRIA group for giving us access to their ranging system, and to Philippe Brou for bringing his ranging system back to life.

References

- Asada, Haruo, and Brady, Michael, [1984], The curvature primal sketch, MIT Artificial Intelligence Laboratory, AIM-758.
- Binford, Thomas O., [1981], "Inferring surfaces from images," *Artificial Intelligence*, 17, 205 — 244.
- Brady, Michael, Ponce, Jean, Yuille, Alan, and Asada, Haruo, [1985], "Describing surfaces," *Comp. Vis. Graph. and Im. Proc.*.
- Brou, P, [1984], "Finding the orientation of objects in vector maps," *Int. J. Rob. Res.*, 3 (4) .
- Burt, Peter J, [1981], "Fast filter transforms for image processing," *Comp. Graph. and Im. Proc.*, 16, 20 — 51.
- Burt, Peter J, [1983], "Fast algorithms for estimating local image properties," *Comp. Graph. and Im. Proc.*, 21, 368 — 382.
- Canny, John Francis, [1983], "Finding Edges and Lines in Images," *MIT Artificial Intelligence Lab. TR-720*.

- Faugeras, O. D., et. al, [1982], Towards a flexible vision system, *Robot vision*, ed. Pugh, Alan, IFS, UK.
- Faugeras, O. D., and Hebert, M., [1983], A 3-D recognition and positioning algorithm using geometrical matching between primitive surfaces, *Proc. Int. Jt. Conf. Artif. Intell.*, Karlsruhe.
- Faugeras, O. D., Hebert, M., Pauchon, E., and Ponce, J., [1984], Object representation, identification, and positioning from range data, *First International Symposium on Robotics Research* Brady, Michael, and Paul Richard (eds). MIT Press, 425 - 446.
- Grimson, W.E.L., and Lozano-Pérez, Tomas, [1984], "Model-Based Recognition and Localization from Sparse Range or Tactile Data," *Int. J. Rob. Res.*, 3 (3), 3 - 35.
- Horn, B. K. P., [1977], "Understanding image intensities," *Artif. Intell.*, 8.
- Ikeuchi, K., Horn, B. K. P., et. al., [1984], Picking up an object from a pile of objects, *First International Symposium on Robotics Research* Brady, Michael, and Paul Richard (eds). MIT Press, 139 - 162.
- Haralick, Robert M., Watson, Layne T., and Laffey, Thomas J., [1983], "The topographic primal sketch," *Int. J. Rob. Res.*, 2 (1), 50 - 72.
- Huffman, D. A., [1971], Impossible objects as nonsense sentences, *Machine Intelligence 6*, Meltzer, B., and Michie, D. (eds.), Edinburgh University Press, 295 - 323.
- Jacobsen, S. C., Wood, J. E., Knutti, D. F., and Biggers K. B., [1984], "The Utah/MIT dextrous hand: work in progress," *Int. J. Rob. Res.*, 3 (4).
- Jacobsen, S. C., Wood, J. E., Knutti, D. F., Biggers, K. B., and Iversen, E. K., [1985], The Version I UTAH/MIT Dextrous Hand, *Proceedings of the Second International Symposium on Robotics Research*, H. Hanafusa and H. Inoue (eds.), MIT Press, Cambridge.
- Kapur, D., Mundy, J. L., Musser, D., and Narendran, P., [1985], Reasoning about three-dimensional space, *IEEE Int. Conf. on Robotics*, St. Louis, Mo., March 1985, 405 - 410.
- Little, J. J., [1985], Recovering shape and determining attitude from Extended Gaussian Images, University of British Columbia, Dept. Computer Science.
- Marr, D., [1976], "Early processing of visual information," *Phil. Trans. Roy. Soc. London*, B275, 843 - 524.
- Marr, D., and Hildreth, E. C., [1980], "Theory of edge detection," *Proc. Roy. Soc. London*, B207, 187 - 217.
- Richter, J., and Ullman, S., [1982], "A model for the spatio-temporal organisation of X and Y type ganglion cells in the primate retina," *Biol. Cyb.*, 43, 127 - 145.
- Salisbury, J. Kenneth, and Craig, John J., [1982], "Articulated hands: force control and kinematic issues," *Int. J. Rob. Res.*, 1 (1), 4 - 17.
- Terzopoulos, D., [1983], "The role of constraints and discontinuities in visible-surface reconstruction," *Proc. 7th Int. Jt. Conf. Artif. Intell.*, Karlsruhe, 1073 - 1077.
- Turner, K., [1974], Computer perception of curved objects using a television camera, Edinburgh University.

- Yuille, A. L., [1983], Zero-crossings on lines of curvature, Proc. Conf. Amer. Assoc. Artif. Intell., Washington, August.
- Yuille, A. L. and Poggio, T., [1983a], "Fingerprints Theorems for Zero-Crossings," *MIT Artificial Intelligence Laboratory AIM-730*.
- Yuille, A.L. and Poggio, T., [1983b], "Scaling theorems for zero crossings," *MIT Artificial Intelligence Laboratory AIM-722*.
- Witkin, A., [1983], "Scale-Space Filtering," *Proc. 7th Int. Jt. Conf. Artif. Intell., Karlsruhe*, 1019 - 1021.

END

FILMED

11-85

DTIC